

Impact of URI Canonicalization on Memento Count

MAT KELLY, Old Dominion University

LULWAH M. ALKWAI, Old Dominion University

MICHAEL L. NELSON, Old Dominion University

MICHELE C. WEIGLE, Old Dominion University

HERBERT VAN DE SOMPEL, Los Alamos National Laboratory

Quantifying the captures of a URI over time is useful for researchers to identify the extent to which a Web page has been archived. Memento TimeMaps provide a format to list mementos (URI-Ms) for captures along with brief metadata, like Memento-Datetime, for each URI-M. However, when some URI-Ms are dereferenced, they simply provide a redirect to a different URI-M (instead of a unique representation at the datetime), often also present in the TimeMap. This infers that confidently obtaining an accurate count quantifying the number of non-forwarding captures for a URI-R is not possible using a TimeMap alone and that the magnitude of a TimeMap is not equivalent to the number of representations it identifies. In this work we discuss this particular phenomena in depth. We also perform a breakdown of the dynamics of counting mementos for a particular URI-R (google.com) and quantify the prevalence of the various canonicalization patterns that exacerbate attempts at counting using only a TimeMap. For google.com we found that 84.9% of the URI-Ms result in an HTTP redirect when dereferenced. We expand on and apply this metric to TimeMaps for seven other URI-Rs of large Web sites and thirteen academic institutions. Using a ratio metric *DI* for the number of URI-Ms without redirects to those requiring a redirect when dereferenced, five of the eight large web sites' and two of the thirteen academic institutions' TimeMaps had a ratio of ratio less than one, indicating that more than half of the URI-Ms in these TimeMaps result in redirects when dereferenced.

1 INTRODUCTION

Memento TimeMaps serve as an index for the mementos for an original resource (URI-R) contained in an archive [16]. Web archives return TimeMaps with a list of identifiers (URI-Ms) for the HTTP transactions observed at archival time. TimeMaps have generally been used as a count of the number of representations of a URI-R present in an archive. But, TimeMaps may include URI-Ms for archived representations (HTTP 2XX), archived redirections (HTTP 3XX), and archived errors (HTTP 4XX or 5XX) [6]. Only the URI-Ms that result in an HTTP 2XX when dereferenced match the general notion of a “capture” of the contents of a webpage. But, the status that results when a URI-M is dereferenced (to the extent of returning an archived entity body) is not present in a TimeMap. Further, TimeMaps do not explicitly return a “count” value to indicate the number of mementos listed in the TimeMap that produce a non-redirecting (non-3XX) HTTP status code when dereferenced. This can cause problems when using the number of URI-Ms in a TimeMap as a proxy for the number of captures of a Web page.

Various tools [1, 8, 9] and access points into the Web archives return a different count of the number of captures for a URI-R depending on the heuristic implemented and the source of the archival listings. For example, the Internet Archive's web interface when queried with the URI-R example.com states, “Saved 11,771 times between January 20, 2002 and May 20, 2016”. The file returned from Internet Archive's CDX endpoint returns 69,162 entries. The TimeMap from Internet Archive for example.com contains 40,641 URI-Ms with a rel value of “memento”. The heuristic of determining how many captures are represented by URI-Ms in a TimeMap cannot be completed without dereferencing.

Researchers can use the inline metadata about the URI-M, without the need to dereference the URI-M in a TimeMap, including its temporal ordering, datetime (through the datetime HTTP Link attribute

[16]), etc. to infer characteristics about a dereferenced memento. However, dereferencing some URI-Ms in a TimeMap produces an HTTP redirect [6] that instructs the client to access a URI-M with a different datetime, to obtain the requested content. For example, a TimeMap for `http://vimeo.com` from Internet Archive contained 199,262 URI-Ms with an associated “rel” value of “memento”. However, when a user accesses over 57% of these URI-Ms, an HTTP Redirect is returned pointing to another memento whose URI-M is in the TimeMap that returns a HTTP Status OK. A different extreme of memento count results when a user accesses the TimeMap from `http://odu.edu`, whose percentage of redirects is around 9.7% of the URI-Ms listed.

Redirection in a Web archive can be attributed to a variety of canonicalization rules including a scheme change (e.g., `http` to `https`), an obsolete subdomain (e.g., `www2` to `www`), a slash added to a URI (`http://foo.com/~joe` to `http://foo.com/~joe/`), among others. Preserving and replaying these redirects allows an archive to accurately reproduce the HTTP transactions that would have occurred when the URI being accessed resided on the live Web.

When a URI-M in a TimeMap is dereferenced, it may redirect to another URI-M listed in the TimeMap. Because of this, the heuristic of counting URI-Ms with relation values of “memento” is an inaccurate means of determining the number of unique representations inferred from a TimeMap. We further emphasize the distinction per the Memento specification that the identifiers for mementos (URI-Ms) in a TimeMap are identifiers for archived HTTP transactions (e.g., transmission of HTTP 2XX, 3XX, 4XX, etc.) rather than identifiers for representations.

Based on the number of URI-Ms in a TimeMap not necessarily resolving to unique mementos when archival redirects are followed, we examined the mementos from contemporarily large TimeMaps to evaluate the patterns and schemes used in Memento canonicalization. Through this, we identify the difference between the number of mementos available as reported by the TimeMap through naive “rel” counting heuristics to the temporally unique mementos identified once these mementos are dereferenced.

2 BACKGROUND

This section includes background information and an overview of the state-of-the-art of archival technologies relevant to this work including Memento aggregation, URI canonicalization, archival indexing formats, and URI-R opacity.

2.1 Memento Aggregation

The Memento Framework [16] allows navigation of Web archives in the dimension of time using content from Web archives and resource versioning systems. A Memento TimeMap (Figure 1) is a structured list of identifiers (URI-Ms) for archived captures (mementos) returned from an archive when queried with a URI-R as the parameter. A TimeMap may also contain references to other TimeMaps and TimeGates.

A Memento aggregator is a software implementation of Memento that takes a URI as the parameter, queries multiple supported archives, combines and temporally orders the returned mementos, and returns this list as a TimeMap to the user. We used the MemGator [1] implementation of a Memento aggregator in collecting data for our analysis. MemGator provides its own heuristic for determining and reporting the number of mementos present in an aggregated TimeMap using the non-standard `X-Memento-Count` HTTP header. Despite this, we used the contents of the aggregated TimeMap returned from MemGator instead of this header as the basis for further investigating the number of mementos present.

2.2 URI Canonicalization

URI canonicalization associates differently formatted URIs [12]. For example, `http://example.com` might be associated with:

```

<http://example.com>; rel="original",
<http://web.archive.org/web/20020120142510/http://example.com/>; rel="
    ↪ first memento"; datetime="Sun, 20 Jan 2002 14:25:10 GMT",
<http://web.archive.org/web/20020804094019/http://www.example.com/>;
    ↪ rel="memento"; datetime="Sun, 04 Aug 2002 09:40:19 GMT",
<http://web.archive.org/web/20160728014649/http://www.example.com/>;
    ↪ rel="memento"; datetime="Thu, 28 Jul 2016 01:46:49 GMT",
<http://web.archive.org/web/20160728114745/http://www.example.com>; rel
    ↪ ="memento"; datetime="Thu, 28 Jul 2016 11:47:45 GMT",
<http://web.archive.org/web/20160728123024/http://example.com/>; rel="
    ↪ last memento"; datetime="Thu, 28 Jul 2016 12:30:24 GMT",
<http://localhost:1208/timemap/link/http://example.com>; anchor="http
    ↪ ://example.com"; rel="timemap"; type="application/link-format",
<http://localhost:1208/timegate/http://example.com>; anchor="http://
    ↪ example.com"; rel="timegate"

```

Fig. 1. A partial Link formatted TimeMap from a local instance of MemGator. Highlighted rel values constitute inclusion in the sum described in Equations 1 and 2.

- <http://www.example.com>
- <https://www.example.com>
- <http://example.com/>
- <http://example.com/index.html>
- <http://example.com/#articles>

Canonicalization allows after-the-fact clustering of URIs that likely reference the same resource. As URI schemes from a Web site change over time, canonicalization is critical for retaining a cohesive, comprehensive listing of the mementos available for a Web page. Internet Archive’s Wayback CDX Server API¹ and endpoint² is one of multiple endpoints³ that provides access to the indexes of the archive’s holdings. A partial example (that corresponds with the TimeMap shown in Figure 1) of the data returned from Internet Archive’s CDX server is shown in Figure 2. As an alternative to their Memento endpoint [11], the CDX endpoint provides the HTTP status code of the capture as well as Sort-friendly URL Reordering Transformed (SURT) URIs. Part of the SURT-generation process involves canonicalizing the URI-R. A canonicalized URI is present in the first space-delimited field (Figure 2) where the “www” subdomain is not present despite being part of the URI in the query parameter. The non-canonicalized URI-R attributed to the record in the CDX is available as the third field in the CDX record. Figure 2 shows the URI-R variations including no subdomain, the “www” subdomain, with and without a trailing slash, and the explicit inclusion of the port number as all canonicalizing to the same URI in the CDX.

2.3 Archival Indexing

A CDX record with a 3XX HTTP status code does not contain the ultimate URI-M that the user will experience when the URI-M is dereferenced. Further, the CDX in Figure 2 is only representative of IA’s

¹<https://github.com/internetarchive/wayback/tree/master/wayback-cdx-server>

²Example access at <http://web.archive.org/cdx/search/cdx?url=example.com>

³https://archive.org/help/wayback_api.php

```

com,example)/ 20020120142510 http://example.com:80/ text/html 200
  ↪ HT2DYGA5UKZCPBSFV3JOBXGW2G5UUA 1792
com,example)/ 20020804094019 http://www.example.com:80/ text/html 200
  ↪ UY3I2DT2AMWAY6DECFCFYMT5Z0TFHUCH 457
com,example)/ 20160728014649 http://www.example.com/ unk 302 3
  ↪ I42H3S6NNFQ2MSVX7XZKYAYSCX5QBYJ 339
com,example)/ 20160728114745 http://www.example.com unk 302 3
  ↪ I42H3S6NNFQ2MSVX7XZKYAYSCX5QBYJ 340
com,example)/ 20160728123024 http://example.com/ text/html 200
  ↪ ASIFPQKKLDWATFDI010JJ3NSK34KLLMN 577

```

Fig. 2. A CDX response returned from Internet Archive’s CDX Server. The space-delimited fields are representative of the canonicalized (SURTed) URI, datetime, original URI, MIME type of original document (where applicable), HTTP response code, digest of WARC response record, and length of response record.

holdings. The corresponding service of aggregating CDX records in an entity like Memento’s concept of combining of TimeMaps from different archives through a Memento aggregator does not exist in standard practice for CDX files. Archives that provide a Memento endpoint are not required and frequently do not expose a CDX endpoint like Internet Archive. This prevents simply referencing all aggregated archives’ CDX files for a URI to determine the non-redirection count of mementos. In this work, we utilize the aggregated holdings of multiple Web archives as well as the CDXJ format [2], an extension of CDX. MemGator’s CDXJ generation is derived from the archives’ Memento endpoints, specifically their Link formatted TimeMaps, and transformed into the CDXJ format that allows quicker, more reliable parsing of the datetime that the included URI-Ms represent.

2.4 Opacity of URI-Ms

It is tempting to extract URI-R and Memento-Datetime values directly from URI-Ms. For example, it is likely that `http://web.archive.org/web/20140417054441/http://google.com/` is a memento for `http://google.com` appeared at April 17, 2014 at 05:44:41 GMT. However, we cannot be sure until we dereference the URI-M and check its response headers for the values in `rel="original"` and Memento-Datetime. While it is unlikely that the IA will deceive us, the URI-M may redirect to another URI-M with a different Memento-Datetime, or in the case of an archived HTTP redirection, the URI-M might end up at an altogether different URI-R [3]. Furthermore, some archives issue URI-Ms without semantics — for example these URI-Ms are all mementos for `google.com` but neither this nor the Memento-Datetime can be ascertained without dereferencing: `webcitation.org/query?id=1398456230796350`, `archive.is/sz8b9`, and `perma.cc/H3YY-BQN5`. For these reasons, we treat URI-Ms as fully opaque [7] and dereference all URI-Ms to extract values for URI-R and Memento-Datetime.

3 RELATED WORK

Bar-Yossef et al. [4] introduced the term, “Soft 404s” to identify Web pages that report a status code other than HTTP 404 despite the page not existing. Meneses et al. [10] described the process of identifying “Soft 404s” based on a signature of the page’s contents. In this work we describe “soft 3XXs” where content is returned from an archive with a status code of 200 yet the contents of the capture consist of

Archive	Memento count
Internet Archive	636,246
Archive-It	62,828
Webcitation	7,551
Stanford Web Archive	4,734
UK National Archives Web Archive	1,510
Archive.is	1,301
PRONI Web Archive	173
UK Parliament Web Archive	127
Total	714,470

Table 1. Distribution of mementos for google.com for a collection of archives defined by a locally deployed Memento aggregator.

an archived HTTP 3XX redirect. With archives that implement Memento, the Accept-Datetime header instructs the archive to return the originally archived status code (Section 6.1).

AlSum et al. [3] analyzed memento redirection patterns relating to HTTP redirects to supply the user with the correct memento when a redirect is encountered in the archives. They introduced the notion of “URI stability” to give a quantitative measure of the presence of HTTP 3XX status codes that result when URI-Ms in TimeMaps are dereferenced.

Rosenthal has discussed Memento aggregator merits and downsides. He cautioned against using TimeMap magnitude for determining the number of URI-Ms available, stating, “Even if we assume that an archive is correct in announcing that it contains a valid copy of the resource at a particular URL at a particular time, that does not imply that it is willing to satisfy a browser’s request for that copy.” [15].

4 DATA COLLECTION

To analyze the degree to which archival identifiers result in redirects, we needed to acquire the HTTP response headers for all URI-Ms accumulated from multiple Web archives for a URI-R. The concept of a Memento aggregator allows us to accomplish this task, albeit parsing the standardized Link formatted resulting TimeMaps is potentially prone to error.

We deployed a local instance of MemGator⁴ version 1.0-RC4 configured to query the archives listed in Table 1. The MemGator instance was initialized with 25 minutes as the value for the “restimeout” (response timeout for each archive) and “hdrtimeout” (header timeout for each archive) parameters. Declaring these timeout values ensured that the server portion of data collection was not prematurely returned because of network latency with communication to the considered archives. Our client script that queried the local MemGator instance was also setup to access this instance with equally large timeout values.

We leveraged MemGator’s CDXJ [2] interface (example output in Figure 3) for simple datetime extraction, structured JSON-formatted metadata of each memento’s attributes, and more human readable output compared to the conventional Link (Figure 1) or JSON formatted TimeMaps. Collection was run on a late 2013 MacBook Pro running OS X version 10.11.4 with a 2.4 GHz Intel i5 processor, 8 GB of RAM, and a 250 GB SSD disk. Data was collected mid-May, 2016. We performed an initial analysis of the mementos contained within the TimeMap without dereferencing any mementos. The client script to

⁴<http://github.com/oduwsdl/memgator>

```

@meta {"original_uri": "http://example.com"}
@meta {"timegate_uri": "http://localhost:1208/timegate/http://example.com"}
@meta {"timemap_uri": {
  "link_format": "http://localhost:1208/timemap/link/http://example.com",
  "json_format": "http://localhost:1208/timemap/json/http://example.com",
  "cdxj_format": "http://localhost:1208/timemap/cdxj/http://example.com"
}}
...
20090418233448 {"uri": "http://web.archive.org/web/20090418233448/http://www.example.com
  ↪ /", "rel": "memento", "datetime": "Sat, 18 Apr 2009 23:34:48 GMT"}
20090421223547 {"uri": "http://wayback.vefsafn.is/wayback/20090421223547/http://www.
  ↪ example.com/", "rel": "memento", "datetime": "Tue, 21 Apr 2009 22:35:47 GMT"}
20090421231335 {"uri": "http://webarchive.loc.gov/all/20090421231335/http://www.example.
  ↪ com/", "rel": "memento", "datetime": "Tue, 21 Apr 2009 23:13:35 GMT"}
...

```

Fig. 3. A partial CDXJ formatted TimeMap returned from a local instance of MemGator containing URI-Ms from multiple archives.

query the local MemGator was created in Python 2.7.10 using the “requests” library⁵ and the built-in JSON parser.

5 ANALYSIS BASED ON TIMEMAPS

We obtained a TimeMap for google.com from our locally deployed Memento aggregator (MemGator instance) containing 714,470 URI-Ms from 8 different Memento-compliant archives. Table 1 shows the distribution of the mementos using a simple URI-based association algorithm to attribute URI-Ms to an archive. 89.1% of the URI-Ms returned were from Internet Archive.

5.1 Variation in Scheme

Two schemes [5] are used for the URI-Rs contained within the URI-Ms returned: HTTP and HTTPS. Table 2 shows the breakdown of the URI-Rs contained in the URI-Ms based on the TimeMap, inclusive of the inferred URI-Rs whose scheme could not be determined solely from the URI-M. As discussed in Section 2.4, the more accurate method to attribute a URI-R to a URI-M is to obtain the memento’s “original” Link header values but this section focuses on analyzing the contents of the TimeMap without requesting the URI-Ms.

From the URI-Rs that could be extracted, 86.2% used the HTTP scheme. Table 3 shows the URI-R-based memento count for each scheme using a substring-based grouping approach similar to that used for Table 1. The TimeMap contained canonicalized variants of the URI-Rs embedded as substrings within the URI-Ms with our query for the CDXJ-formatted TimeMap. Our query supplied the URI-R variant containing the HTTP scheme, www sub-domain, and no trailing characters combination (i.e., http://www.google.com) via the query to MemGator⁶.

5.2 Grouping by Year

We separated the URI-Ms as reported by the TimeMap into year-based buckets using the “datetime” attribute for each URI-M (and not the embedded 14-digit date stamp per Section 2.4) as well as by-archive

⁵<http://python-requests.org>

⁶<http://localhost:1208/timemap/cdxj/http://www.google.com>

Scheme	URI-R count
http	609,274
https	97,645
unknown	7,551
	714,470

Table 2. Scheme distribution among the URI-Rs within the mementos in the TimeMap for google.com.

Scheme	Format	URI-R count
http	http://www.google.com	541,160
	http://google.com	67,811
	http://other.google.com	303
https	https://www.google.com	96,853
	https://google.com	792
	https://other.google.com	0

Table 3. Count of URI-Rs contained within URI-Ms for google.com.

for google.com. We calculated the average time between URI-Ms within a year-based bucket to show that the velocity of capturing google.com is generally increasing in time (Figure 4). The quantity of captures from Internet Archive for google.com in 2015 was significantly lower than the trend would indicate. Table 4 indicates this dramatic drop in google.com captures from both the IA CDX endpoint and from the Memento endpoint. Also because data collection occurred in May 2016, the partial year data points for 2016 are on-par with the trend of years prior to 2015.

5.3 “TimeMap” from CDX Server

We also obtained the CDX for google.com from Internet Archive (IA). We compared the HTTP response codes we received when dereferencing the IA URI-Ms from the CDXJ TimeMap (Section 4) with the response codes explicitly provided in the CDX listing IA returned. Per Section 2.2, a CDX endpoint is not available as a user-accessible endpoint from most Web archives. We used the available endpoint at IA as a sanity check for correctness of the data obtained when the URI-Ms in a TimeMap are dereferenced. The intention of analyzing the TimeMap and not simply deferring to the CDX Server, despite the majority of mementos in the aggregated TimeMap being from IA, is to extrapolate the dereferencing strategy to other Memento-compliant Web archives.




Fig. 4. The average time between consecutive mementos has decreased with time. This plot shows a year-based bucketing of the difference in time between adjacent mementos from different archives for google.com.

year	M_{TM}	M_{RC}	DI
1998	4	4	∞
1999	19	19	∞
2000	132	87	1.933
2001	1,185	579	0.955
2002	176	137	3.513
2003	75	55	2.750
2004	197	143	2.648
2005	1,236	414	0.504
2006	735	483	1.917
2007	1,055	842	3.953
2008	1,376	894	1.855
2009	6,074	4,335	2.493
2010	9,326	6,530	2.335
2011	20,634	9,279	0.817
2012	102,533	16,240	0.188
2013	228,405	25,203	0.124
2014	164,865	22,738	0.160
2015	17,978	11,286	1.686
2016	139,520	5,805	0.043

Table 4. Google over time, bucketed by year, based on IA mementos extracted from the MemGator CDXJ TimeMap. M_{TM} is the memento count based solely on the data in the TimeMap, M_{RC} is the count based on exclusion of redirects when dereferenced, and DI is the ratio of non-redirecting mementos to redirecting mementos, per Section 6.3.

6 ANALYSIS BASED ON MEMENTOS

In Section 5 we analyzed the archival presence of a URI based solely on the TimeMap supplied by a local aggregator when querying the aggregator with one canonicalized variant of the URI-R. In this section, we dereference the URI-Ms in the TimeMap for further analysis. Equations 1 and 2 set the basis for counting mementos in a TimeMap by merely counting the entries where the “rel” attribute in a TimeMap contains a value of “memento”. For example, Figure 1 shows a Link formatted TimeMap where each highlighted entry containing a “memento” rel value (with the potential additional inclusion of other values like “last” and “first”) increments the count according to Equation 2.

$$M = \text{URI-M} \subset \text{TimeMap, if } \text{“memento”} \subset \text{valuesOf}(\text{rel}) \quad (1)$$

$$|TM|_{rel} = \sum_{m=1}^{lenM} 1 \quad (2)$$

6.1 Redirects in Mementos

The number of non-redirecting (non-3XX) mementos in a TimeMap cannot be counted with the TimeMap data alone. When URI-Ms are dereferenced, they need not contain an entity body but may consist only of an archived HTTP response, which might not be a 200. This occurs in cases where the live Web site returned an HTTP 302 redirect, among other circumstances. This redirect was captured, retained, and is replayed by the archives. When replayed, the datetime originally requested for a URI-R, as is often the case, will be different than the datetime of the memento ultimately served to the user. This “archived 302” is different from a 3XX code returned from an archive that is not representative of an archival capture; for example, when a datetime for a URI is requested where no capture for the URI is contained within the archive’s holdings.

Assuming the TimeMap in Figure 1 is wholly inclusive of all of the mementos contained by Internet Archive for example.com, requesting the URI-M <http://web.archive.org/web/20160728114743/http://www.example.com> (two seconds before a listed URI-M neglecting datetime semantics per Section 2) will result in a 302 from the archive pointing to the nearest capture. This behavior is a function of the archive, is not mandatory behavior to exhibit, and is functionality independent of the Memento protocol. Were there a capture at the former datetime where the archival crawler experienced a 302 from the live Web at the time, the TimeMap would contain the URI-M <http://web.archive.org/web/20160728114743/http://www.example.com> with a rel value of “memento” indistinguishable from the memento entry at <http://web.archive.org/web/20160728114743/http://www.example.com> regardless of the status code that occurs from the archive when each URI-M is dereferenced.

6.2 Direct and Indirect Mementos

Users interacting with an archive via a Web browser will not directly experience intermediary HTTP transactions (the user agent automatically redirects the user to a non-3XX status), we introduce the term URI-M_D (for “direct”) to indicate a URI-M in a TimeMap that does not require any intermediary transaction to resolve. Thus, a URI-M_D is a case of a URI-M where the URI-M originally requested by the user is the identifier for the ultimate memento served. URI-Ms in a TimeMap that exhibit the behavior where a HTTP 3XX class code is replayed and the datetime differs from that requested by the user are indicated with URI-M_I (for “indirect”).

In Equation 3 we filter $|TM|_{rel}$ from Equation 2 to exclude mementos that resolve to HTTP 3XX status codes. $|TM|_D$ represents the count of mementos that result in non-3XX statuses based on the

URI-Ms in a TimeMap. Section 2.2.4 of the Memento RFC [16] states that a link with a datetime attribute must match the value of the Memento-Datetime header when the link is dereferenced.

$$|TM|_D = \sum_{m=1}^{lenM} \begin{cases} 0 & 300 \geq httpStatus_m < 400, \\ 1 & \text{otherwise.} \end{cases} \quad (3)$$

$$|TM|_I = |TM|_{rel} - |TM|_D \quad (4)$$

We quantify the ratio of mementos with non-redirecting HTTP status codes (Equation 3) to those with redirects (Equation 4) in Equation 5 as DI .

$$DI = \begin{cases} \frac{|TM|_D}{|TM|_I} & |TM|_I > 0, \\ \infty & \text{otherwise.} \end{cases} \quad (5)$$

As an example, Figure 5 contains 11 URI-Ms that result in non-redirecting archived HTTP status codes when dereferenced, inclusive of eight 200 codes, two 4XX codes, and one archived 504. TimeMap A represents a domain for an organization (for example) that is acquired by another organization, whose domain is represented by TimeMap B. At the point of acquisition, TimeMap A redirects to TimeMap B, as represented by the HTTP 301. At two points prior to acquisition, an archival crawler attempted to capture the URI-R for TimeMap A but received a server-side redirect, which is reflected in the preserved HTTP 302 responses. As the acquisition proceeded, the URI-R may have been deleted (the 404 in TimeMap A) and the server misconfigured in the transition (HTTP 504). An intermittent HTTP 401 (Unauthorized) error is also experienced in the URI-R for TimeMap B as the servers are reconfigured to accept the additional traffic from the acquisition.

Three URI-Ms resulted in 3XX redirects when dereferenced. Using Equation 5, $DI = 113 \approx 3.7$. Sparsely archived URIs will often contain a list of URI-Ms where all result in an HTTP 200 status code when dereferenced, which would result in DI being undefined. It is far less likely that all URI-Ms in a TimeMap return in an HTTP redirect when dereferenced.

During the data acquisition process (Section 4) we experienced intermittent HTTP 5XX status codes [6] in responses from Internet Archive, namely HTTP 503 (Service Unavailable) and 504 (Gateway Timeout). In much of the same way that sending an Accept-Datetime header causes a “soft” HTTP status code to “harden”, we repeated the request via curl⁷ with the inclusion of an Accept-Datetime HTTP header [16]. This additional step caused no change in the subsequently returned results compared to the originally results response. Repeating the collection procedure for select URI-Ms in the future remedied this issue, allowing us to attribute the error to the archive and not the archive returning a capture of an archived 5XX. If the response instead indicated that the returned 5XX status codes was representative of the state of the URI-R at the respective time (through providing a Memento-Datetime response header) and not an intermittent result attributable to the archive, the URI-M would signify an increment in the Equation 3 summation.

6.3 Canonicalization Patterns

In observing the mementos for <http://www.google.com>, we encountered 3 canonicalization patterns for URI-Ms that surface those that are $URI-M_D$ versus those that are $URI-M_I$. We define M_{TM} to be the memento count for a TimeMap when using only the data contained in the TimeMap without dereferencing mementos (Section 5). We define the *representative count* (M_{RC}) of the number of mementos present in a

⁷<https://curl.haxx.se/>

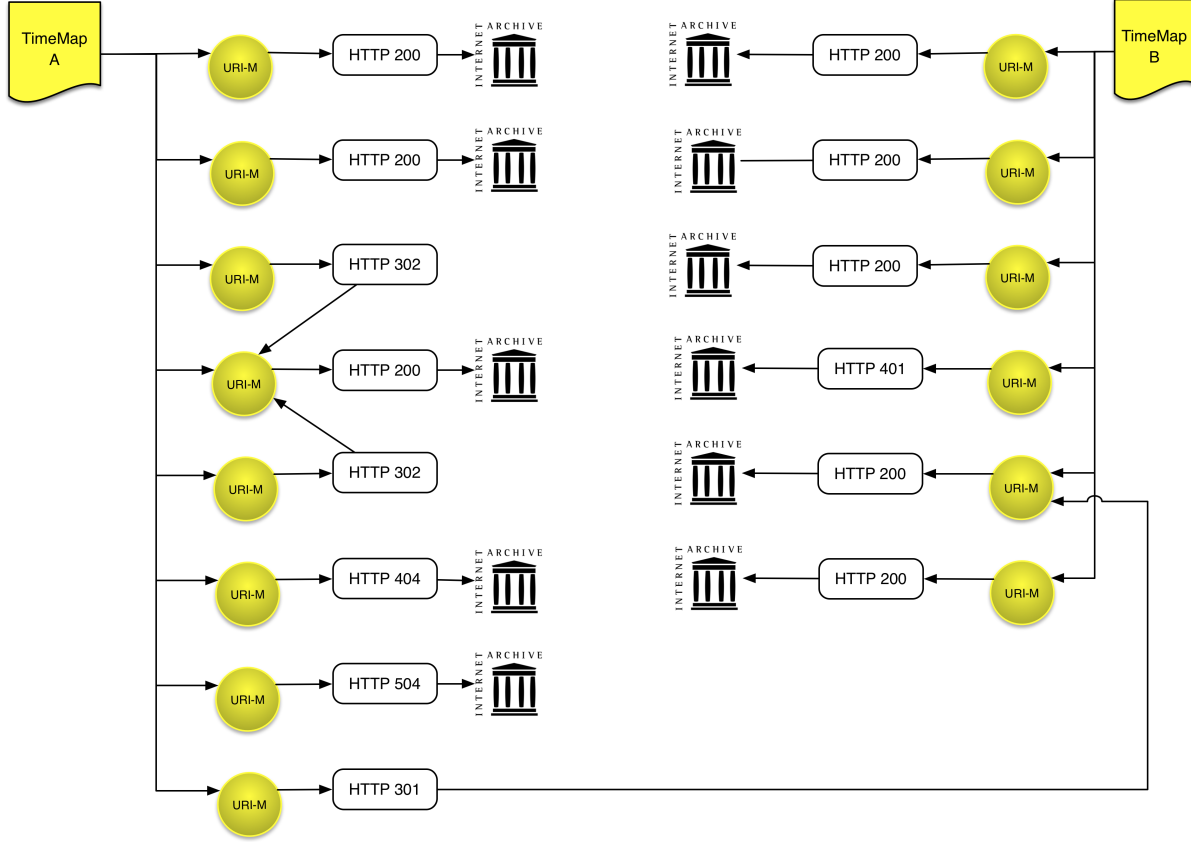


Fig. 5. Dereferencing URI-Ms in a TimeMap may produce a variety of HTTP status code, some of which are redirects both to other URI-Ms within the TimeMap and URI-Ms not included in the initial TimeMap. Counting the number of mementos without dereferencing URI-Ms is therefore problematic.

TimeMap to be the number of URI-M_D where $M_{RC} \leq M_{TM}$. These canonicalization patterns observed are Inter-scheme, Slash-added, and subdomain redirect patterns, described in this section.

6.3.1 Inter-scheme URI-M Redirect. As adoption of the secure HTTPS scheme over HTTP becomes more prevalent on the live Web [13, 14], the trend becomes apparent in the archives through canonicalizing the HTTP and HTTPS site to be one in the same. For example, observe two mementos from 2013 from the TimeMap for google.com (Figure 6). The status code returned for M_2 is 200 with no HTTP location header present (an example of a URI-M_D). However, the status code returned for M_3 is an HTTP 302 with an HTTP location response header of `/web/20130101000813/http://www.google.com/`, i.e., M_3 redirects to M_2 when dereferenced. Thus, M_3 is a URI-M_I . Were the naive but often applied Equation 2 used for determining how many mementos are represented by the URI-Ms M_2 and M_3 , both would be included while dereferencing each URI-M would result in a count of only a single memento. This highlights an important distinction and discrepancy between the number of identifiers (URI-Ms) and the number of representations (mementos).

```

<http://google.com>; rel="original",
...
M1 <http://web.archive.org/web/20011124163711/http://www2.google.com/>; rel="memento";
    ↪ datetime="Sat, 24 Nov 2001 16:37:11 GMT",
...
M2 <http://web.archive.org/web/20130101000813/http://www.google.com/>; rel="memento";
    ↪ datetime="Tue, 01 Jan 2013 00:08:13 GMT",
M3 <http://web.archive.org/web/20130101003310/https://www.google.com/>; rel="memento";
    ↪ datetime="Tue, 01 Jan 2013 00:33:10 GMT",
...
M4 <http://web.archive.org/web/20140425221431/http://www.google.com>; rel="memento";
    ↪ datetime="Fri, 25 Apr 2014 22:14:31 GMT",
M5 <http://web.archive.org/web/20140425221433/https://www.google.com/>; rel="memento";
    ↪ datetime="Fri, 25 Apr 2014 22:14:33 GMT",
...
M6 <http://web.archive.org/web/20160519223823/http://www.google.com/>; rel="memento";
    ↪ datetime="Thu, 19 May 2016 22:38:23 GMT",
M7 <http://web.archive.org/web/20160520165954/http://google.com/>; rel="memento"; datetime=
    ↪ "Fri, 20 May 2016 16:59:54 GMT",
...

```

Fig. 6. A partial TimeMap in Link format for google.com with annotations highlighting various URI-Ms, discussed in Section 6.3.

6.3.2 Slash-added URI-M Redirect. Queries for the URI `http://www.google.com` are sometimes redirected to the same URI with an appended slash. For example, the two mementos M_4 and M_5 both exist in a TimeMap (Figure 6).

When dereferenced, M_4 returns a 302 with a location header pointing to M_5 , captured two seconds later based solely on the embedded datetime. Both URI-Ms are reported by the TimeMap while only the latter contains an entity body when dereferenced.

6.3.3 Subdomain URI-M Redirect. It is also useful to observe canonicalization that does not result in a redirect. Google has used a variety of subdomains of the sort containing the literal “www” followed by a digit over the years, as with M_1 (Figure 6). Accessing this memento (dereferencing M_1) results in an HTTP 200 status code. The other, much more common subdomain of www, as with M_6 returns an HTTP 302 redirected to M_7 , both present in the TimeMap. Table 7a shows the magnitude of redirects for google.com based on the URI- R_{orig} scheme, URI- R_{orig} subdomain, URI- R_{dest} scheme, and URI- R_{orig} subdomain. The breakdown in Table 7a was also generated for comparison to the URI-Rs vimeo.com and wikipedia.org in Tables 7b and 7c, respectively.

6.3.4 Analysis. Juxtaposing the proportion exhibited for each of the four permutations of scheme transitions (HTTP-to-HTTP, HTTP-to-HTTPS, etc.) from vimeo.com and wikipedia.org as compared to google.com, the inter-scheme transition seems more common with the former pair while the bulk of the results for google.com reside in redirects that retain both the HTTP scheme in the URI-R but also the www subdomain. Focusing specifically on the HTTP-to-HTTPS inter-scheme transition, Figure 7 serves as an interesting cross-section of Table 7a broken down by time. Disregarding the anomalous captures from Internet Archive in 2015, the overall trend of inter-scheme redirects is leaning toward more HTTP-to-HTTPS than HTTPS-to-HTTP as the secure scheme is adopted by more sites on the live Web. Disregarding 2015, the HTTP-to-HTTPS redirects for vimeo.com appear to be monotonically increasing

Time Gap Bucket	URI-M count
0 seconds	24,541
1 second	34,577
2 seconds	26,153
3 seconds	62,526
4 seconds	46,738
5 seconds	14,215
6 seconds	12,213
7 seconds	9,748
8 seconds	7,431
9 seconds	6,610
> 9 seconds, \leq 1 minute	101,868
> 1 minute, \leq 1 hour	247,192
> 1 hour, \leq 1 day	37,399
> 1 day	5,034

Table 5. A range of time differences exists between adjacent captures of a URI-R. This table represents the instances of these differences between adjacent URI-Ms from the TimeMap for google.com reported by Internet Archive.

with normalization for the partial year results for 2016 (collection was performed in May of that year). Figure 7 also shows a steeper quantity of captures containing these redirects for wikipedia.org with few captures of redirects exhibiting this inter-scheme permutation prior to 2014. The rapid increase in each site may temporally correspond with the adoption of the HTTPS scheme by the live Web site, thereby forwarding all traffic accessing the HTTP version of the site using an HTTP 3XX response.

An additional nuance to account for the large quantity of redirects from HTTP URI-Ms to HTTP URI-Ms for google.com can be observed by the large quantity of “revisit” entries in IA’s CDX results for google.com. A revisit entry occurs when an archival crawler is returned content that is identical to a previous capture, often attributed using a hashing scheme on the live page’s content. If an archive reports revisit records as an HTTP redirect based on the CDX listing, and this redirect is propagated to the archive’s Memento endpoint thus producing a unique URI-M, the *DI*’s value for the URI-R decreases. Requesting the URI-M using the Accept-Datetime HTTP header then observing the Memento-Datetime response header’s presence often reveals this nuance, but by relying on the TimeData data without requesting each URI-M, the *DI* for the URI-R is unknown.

6.4 Inter-Memento Temporality

We measured the time between each pair of consecutive mementos, shown in Table 5. We found that 38.4% had a time gap less than 9 seconds, indicating (in some cases) that a redirect would occur when the URI-M is dereferenced. Using a yearly bucketing scheme (Table 10), we plotted the difference in time between adjacent mementos from IA based on the scheme in the URI-M_I and the ultimate URI-M that results when the URI-M_I is dereferenced. For each log-log plot, shown with more details in the Appendix, a point’s quadrant positioning is indicative of the quantity of mementos with a seconds-level granularity of time. For example, a point in the top-left quadrant of a plot indicates that there are many temporally consecutive memento pairs with a very small time difference between them. Top right would indicate many pairs with a large time difference between them; bottom right: few memento pairs with a large time difference; bottom left: few memento pairs with a small time difference. Many more points in the

Year	Memento count
2001	68
2005	391
2006	8
2007	81
2008	62
2009	153
2010	124
2011	616
2012	5,564
2013	25,914
2014	40,819
2015	1,367
2016	10,104

Table 6. Google memento analysis group by memento year where the time between two mementos is less than or equal to 2 seconds.

left half of a plot than the right indicates much less time between captures, i.e., the capture frequency was higher that year. More points being in the right half of the plot indicates that more time passes between consecutive captures. The trend for google.com excluding 2015 shows fewer pairs with a small time difference (more points in the bottom right) as time goes on for all redirect patterns other than HTTP-to-HTTP.

6.5 Temporal Closeness as an Indicator of Redirection

In only brief examination of the TimeMap, some temporally consecutive URI-Ms appeared in “pairs” where a second URI-M exists from the same archive within seconds of the previous. Table 6 lists the by-year breakdown filtering to only include the URI-Ms where the time between the two is less than two seconds. The trend generally increases with time. This plot can also be cross-referenced with Table 7a, which shows the overall inter-scheme redirect breakdown totals independent of time with the additional subdomain granularity.

Table 6 also shows a peak in 2014 at 40,819 pairs where the redirect is less than or equal to two seconds apart from the URI- R_{orig} to the URI- R_{dest} . Given the quantity of inter-scheme redirects in Figure 7 for 2014 totaling around 30,000 as the sum and Table 7a showing a significantly larger number of same scheme redirects (e.g., 490,836 just for HTTP-to-HTTP both with the www subdomain), many redirects over the archived history of google.com can be attributed to something other than a scheme switch. The large number of aforementioned identical scheme and subdomain redirects indicates patterned responses like slash-added (Section 6.3.2) rather than scheme (Section 6.3.1) or subdomain switch (Section 6.3.3).

6.6 Beyond Google

We then evaluated the applicability of the observations for google.com with other archived Web sites. We dereferenced the TimeMaps of 7 additional large Web sites (Table 8) with a variety of adoption trends of HTTPS and ephemerality as well as 13 home pages of various universities and colleges (Table 9). From this further analysis, we observed how prevalent the trend is as exhibited by google.com with a hypothesis




Fig. 7. Inter-scheme redirects for google.com, vimeo.com, and wikipedia.org from the mementos in the TimeMap from MemGator for IA that result in a 3XX.

that the relatively static, fundamentally unchanging Google homepage is a reason for the relatively low DI .

For the select academic institutions in Table 9, DI is inversely proportion to M_{TM} , albeit not strictly as evidenced by “gatech” and “odu”. This pattern does not generally hold in comparison to the large

URI- R_{orig} scheme \ URI- R_{dest} scheme		http			https		
		none	www	other	none	www	other
http	none	1,279	68,837	55	12	20,825	27
	www	8,934	490,836	204	32	77,610	16
	other	0	224	22	0	26	2
https	none	14	731	0	0	296	1
	www	1,117	72,874	27	15	18,525	2,101
	other	0	0	0	0	0	0

(a) Scheme and subdomain for redirects when dereferencing URI-Ms for google.com.

URI- R_{orig} scheme \ URI- R_{dest} scheme		http			https		
		none	www	other	none	www	other
http	none	1,642	104	0	82,637	0	0
	www	1,273	50	0	6,355	0	0
	other	0	0	0	0	0	0
https	none	315	6	0	35,293	1	0
	www	10	0	0	1,149	0	0
	other	0	0	0	0	0	0

(b) Scheme and subdomain for redirects when dereferencing URI-Ms for vimeo.com.

URI- R_{orig} scheme \ URI- R_{dest} scheme		http			https		
		none	www	other	none	www	other
http	none	91	10,575	0	0	4,140	0
	www	110	5,099	0	1	44,104	0
	other	0	0	0	0	0	0
https	none	1	46	0	1	804	0
	www	14	1,014	0	0	5,602	0
	other	0	0	0	0	0	0

(c) Scheme and subdomain for redirects when dereferencing URI-Ms for wikipedia.org.

Table 7. When URI-Ms for three select domains (google.com, vimeo.com, and wikipedia.org), are dereferenced and produce an HTTP redirect, the originally accessed URI- R_{orig} can result in a URI- R_{dest} with a different scheme and subdomain. Cell colors correspond to lines in Figure 7. The scheme and subdomain is “Unknown” for URI-Ms (like those from webcitation) that obfuscate URI-R with which the URI-M is associated. See Section 2.

sites in Table 8 though the selection of sites for each may contain some inadvertent bias. Figure 8 shows nine plots representing the percentage of redirects over time as determined when all URI-Ms with a rel value in the respective TimeMaps from IA are dereferenced.

host	% 3XX	% 200	M_{TM}	DI
google	84.89	15.11	695,525	0.178
yahoo	88.16	11.83	418,896	0.134
sourceforge	73.34	26.63	31,408	0.363
instagram	67.32	32.65	55,228	0.485
vimeo	57.04	42.94	199,262	0.752
cnm	49.97	50.01	87,148	1.001
wikipedia	44.62	55.19	25,973	1.240
whitehouse	44.57	55.24	26,006	1.243

Table 8. Dereferencing 7 other large Web sites' TimeMaps from Internet Archive produces the above distribution of status codes for each site.

host	% 3XX	% 200	M_{TM}	DI
stanford	62.14	37.84	19,309	0.609
princeton	60.10	39.88	9,355	0.663
columbia	48.01	51.88	9,882	1.082
harvard	33.91	65.96	7,699	1.948
caltech	33.13	66.86	5,474	2.017
mit	26.57	73.24	6,379	2.763
gatech	26.03	73.94	3,907	2.841
ufl	24.76	75.23	4,927	3.038
vt	23.07	76.92	4,061	3.334
lsu	15.06	84.93	2,974	5.638
nsu	13.82	86.00	1,208	6.233
odu	9.727	90.27	1,727	9.279
tcc	5.429	94.57	884	17.41

Table 9. Dereferencing the TimeMaps from 13 academic institutions' Web sites from Internet Archive produces the above distribution of status codes for each site.

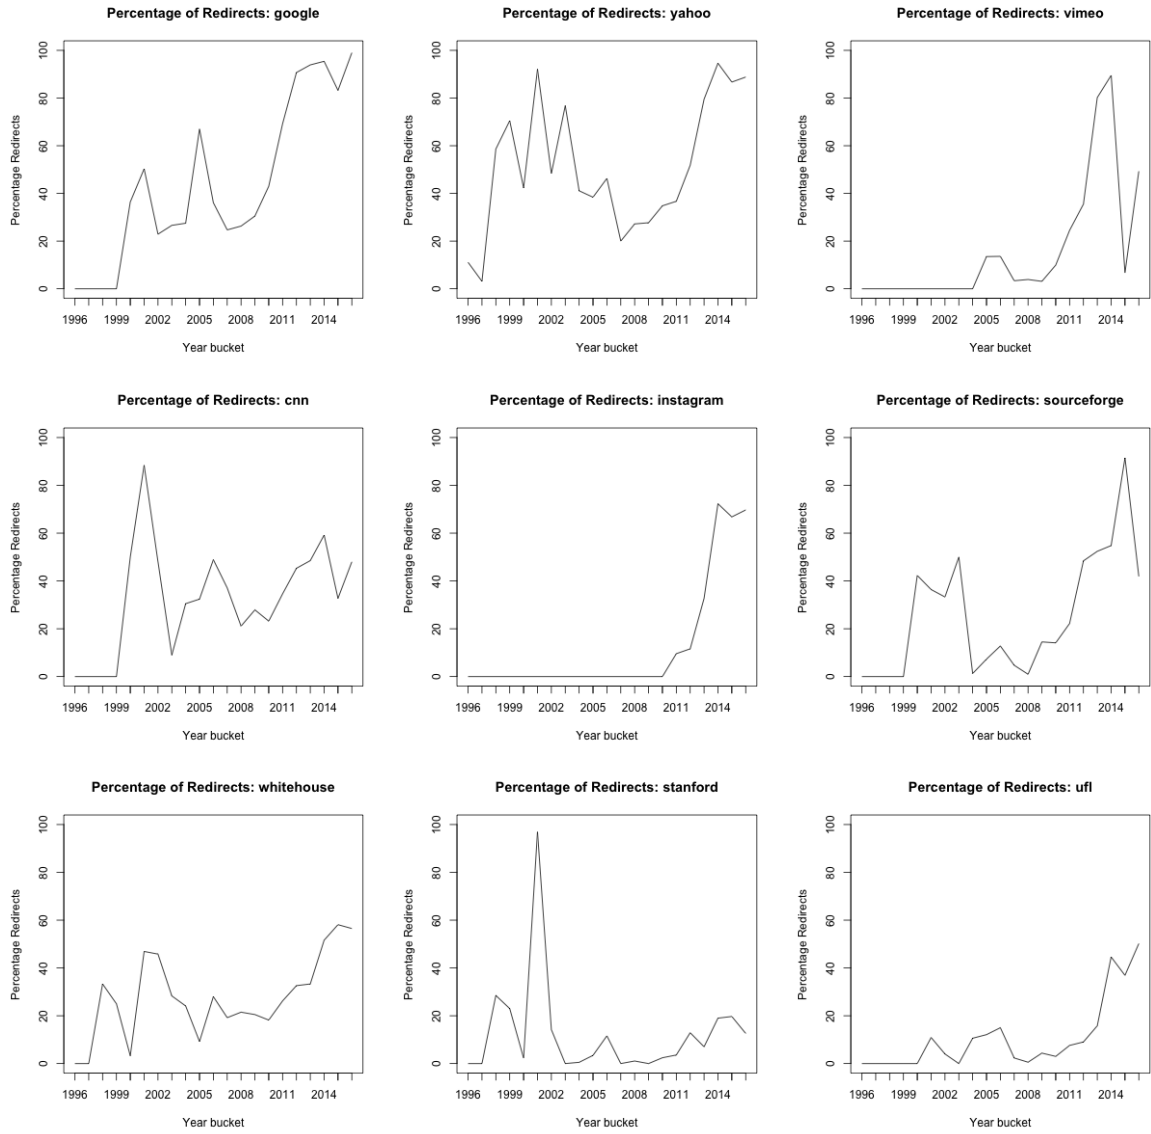


Fig. 8. Nine URI-Rs from Tables 8 and 9 exhibit different degrees of redirection over time.




Table 10. Count of time delta instances for google.com from IA TimeMaps with 3XX redirects. Highlighted columns correspond to Figures 7a and 7. The plots in this Figure are available with more detail in the Appendix.

7 CONCLUSIONS

In this work we identified the problem of attempting to count the number of mementos in a TimeMap based solely on the contents of the TimeMap. We progressively built a method for counting the number of archived captures of Web pages that contain content when dereferenced from a TimeMap. Through observing google.com, a URI-R with a contemporarily large apparent number of mementos, we dereferenced all URI-Ms in an aggregated TimeMap for the URI-R to show that a 84.9% of the URI-Ms are redirects to other URI-Ms in the TimeMap.

We establish the nomenclature of M_{RC} , a means of communicating the number of URI-Ms in a TimeMap that contain a capture with an entity body, as compared to the more naive M_{TM} as calculated using solely the contents of the TimeMap for a URI-R. We analyzed the TimeMaps for the URI-Rs of seven other contemporary large web sites and the TimeMaps from 13 academic institutions. We introduced the DI metric to evaluate the ratio of non-redirecting URI-Ms in a TimeMap to the ratio of redirecting URI-Ms when all URI-Ms in a TimeMap are dereferenced. Five of the eight large Web sites' URI-Rs and two of the thirteen academic institutions' URI-Rs contained more redirecting than non-redirecting mementos when dereferenced ($DI < 1.0$).


From the URI-Ms for google.com that redirected, we split the results between those that changed schemes and those that maintained the same URI-R scheme after the redirect. We split the results on an annual basis to show the effect that the introduction of HTTPS has had on URI-R canonicalization over time. We found that despite an anomalous set of captures in the year 2015, the number of redirects per year on the live Web from HTTP URI-Rs to HTTPS URI-Rs as preserved by the archive has superseded the number of redirects of HTTPS URI-Rs to HTTP URI-Rs. Though the quantity of holdings by Internet Archive for redirects from HTTPS to HTTP is not yet larger than the total of other permutations of HTTP(S) to HTTP(S) redirects (Table 7a), the rapid growth of redirects to the secure scheme for captures confirms and quantifies the increased adoption of HTTPS on the live Web.

REFERENCES


- [1] Sawood Alam and Michael L. Nelson. 2016. MemGator - A Portable Concurrent Memento Aggregator. In *Proceedings of the 16th ACM/IEEE-CS on Joint Conference on Digital Libraries (JCDL)*. 243–244. DOI:http://dx.doi.org/10.1145/2910896.2925452
- [2] Sawood Alam, Michael L. Nelson, Herbert Van de Sompel, Lyudmila Balakireva, Harihar Shankar, and David S. H. Rosenthal. 2015. Web Archive Profiling Through CDX Summarization. In *Proceedings of Theory and Practice of Digital Libraries (TPDL)*. 3–14. DOI:http://dx.doi.org/10.1007/s00799-016-0184-4
- [3] Ahmed AlSum, Robert Sanderson, Herbert Van de Sompel, and Michael L. Nelson. 2013. Archival HTTP Redirection Retrieval Policies. In *Proceedings of the Third Temporal Web Analytics Workshop*. DOI:http://dx.doi.org/10.1145/2487788.2488117
- [4] Ziv Bar-Yossef, Andrei Z. Broder, Ravi Kumar, and Andrew Tompkins. 2004. Sic Transit Gloria Telae: Towards an Understanding of the Web's Decay. In *Proceedings of the 13th International Conference on World Wide Web (WWW)*. 328–337. DOI:http://dx.doi.org/10.1145/988672.988716
- [5] R. Fielding and J. Reschke. 2014. Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing, Internet RFC-7230. (2014).
- [6] R. Fielding and J. Reschke. 2014. Hypertext Transfer Protocol (HTTP/1.1): Semantics and Content, Internet RFC-7231. (2014).
- [7] Ian Jacobs and Norman Walsh. 2004. Web Architecture : URI Opacity. <https://www.w3.org/TR/webarch/#uri-opacity>. (2004). <https://www.w3.org/TR/webarch/#uri-opacity>
- [8] Wesley Jordan, Mat Kelly, Justin F. Brunelle, Laura Vobrak, Michele C. Weigle, and Michael L. Nelson. 2015. Mobile Mink: Merging Mobile and Desktop Archived Webs . In *Proceedings of the ACM/IEEE Joint Conference on Digital Libraries (JCDL)*. 243–244. DOI:http://dx.doi.org/10.1145/2756406.2756956
- [9] Mat Kelly, Michael L. Nelson, and Michele C. Weigle. 2014. Mink: Integrating the Live and Archived Web Viewing Experience Using Web Browsers and Memento. In *Proceedings of the IEEE/ACM Joint Conference on Digital Libraries*

- (JCDL). 469–470. DOI:<http://dx.doi.org/10.1109/JCDL.2014.6970229>
- [10] Luis Meneses, Richard Furuta, and Frank Shipman. 2012. Identifying “Soft 404” Error Pages: Analyzing the Lexical Signatures of Documents in Distributed Collections. In *Proceedings of the International Conference on Theory and Practice of Digital Libraries (TPDL)*. 197–208. DOI:http://dx.doi.org/10.1007/978-3-642-33290-6_22
 - [11] Michael L. Nelson. 2013. Wayback Machine Upgrades Memento Support. <http://ws-dl.blogspot.com/2013/07/2013-07-15-wayback-machine-upgrades.html>. (September 2013).
 - [12] M. Ohye and J. Kupke. 2012. The Canonical Link Relation, Internet RFC-6596. (2012).
 - [13] Guy Podjarny. 2016. HTTPS Adoption *doubled* this year | Snyk. (July 2016). <https://snyk.io/blog/https-breaking-through/> [Online; accessed 21-January-2017].
 - [14] Garrett Robinson and Trevor Timm. 2015. Introducing Secure The News, an automated tool tracking the adoption of HTTPS encryption across news websites. (December 2015). <https://freedom.press/news/introducing-secure-news-automated-tool-tracking-adoption-https-encryption-across-news-websites/> [Online; accessed 21-January-2017].
 - [15] David S. H. Rosenthal. 2013. Re-thinking Memento Aggregation. <http://blog.dshr.org/2013/03/re-thinking-memento-aggregation.html>. (March 2013).
 - [16] Herbert Van de Sompel, Michael Nelson, and Robert Sanderson. 2013. HTTP Framework for Time-Based Access to Resource States – Memento. IETF RFC 7089. (December 2013).


A DETAILED PLOTS FROM TABLE 10




Count of time delta instances for Google.com from archive.org TimeMap
with 3XX redirect from http to http (2011)




Count of time delta instances for Google.com from archive.org TimeMap
with 3XX redirect from https to http (2011)




Count of time delta instances for Google.com from archive.org TimeMap
with 3XX redirect from http to http (2012)




Count of time delta instances for Google.com from archive.org TimeMap
with 3XX redirect from http to https (2012)




Count of time delta instances for Google.com from archive.org TimeMap
with 3XX redirect from https to http (2012)




Count of time delta instances for Google.com from archive.org TimeMap
with 3XX redirect from https to https (2012)




Count of time delta instances for Google.com from archive.org TimeMap
with 3XX redirect from http to http (2013)




Count of time delta instances for Google.com from archive.org TimeMap
with 3XX redirect from http to https (2013)




Count of time delta instances for Google.com from archive.org TimeMap
with 3XX redirect from https to http (2013)




Count of time delta instances for Google.com from archive.org TimeMap
with 3XX redirect from https to https (2013)




Count of time delta instances for Google.com from archive.org TimeMap
with 3XX redirect from http to http (2014)




Count of time delta instances for Google.com from archive.org TimeMap
with 3XX redirect from http to https (2014)




Count of time delta instances for Google.com from archive.org TimeMap
with 3XX redirect from https to http (2014)




Count of time delta instances for Google.com from archive.org TimeMap
with 3XX redirect from https to https (2014)




Count of time delta instances for Google.com from archive.org TimeMap
with 3XX redirect from http to http (2015)




Count of time delta instances for Google.com from archive.org TimeMap
with 3XX redirect from https to http (2015)




Count of time delta instances for Google.com from archive.org TimeMap
with 3XX redirect from http to http (2016)



Count of time delta instances for Google.com from archive.org TimeMap
with 3XX redirect from http to https (2016)



Count of time delta instances for Google.com from archive.org TimeMap
with 3XX redirect from https to http (2016)



Count of time delta instances for Google.com from archive.org TimeMap
with 3XX redirect from https to https (2016)

